



PRESENTS

NETWORLD INTEROP




Content Delivery Networks: The strategy of business-enabling delivery systems

Alistair A. Croll

CEO, Coradiant inc.

3:15 to 4:15, Tuesday, September 11, 2001



Agenda

- The content problem
- Content delivery networks to the rescue
 - How they work
 - Emerging technologies
- Migrating to CDN solutions
 - Separating theory from practice

The content problem

- All content is not alike
 - Some static, some dynamic, some one-time
 - Some confidential, some public
 - Some of broad interest, some narrowcast
 - Some time-sensitive, some robust
- Getting content takes time
 - Speed of light, congestion, protocol overhead

Content delivery networks

- Put copies of content at the network edge
 - Also put some processing at the edge
- Direct clients to the “best” fulfillment point

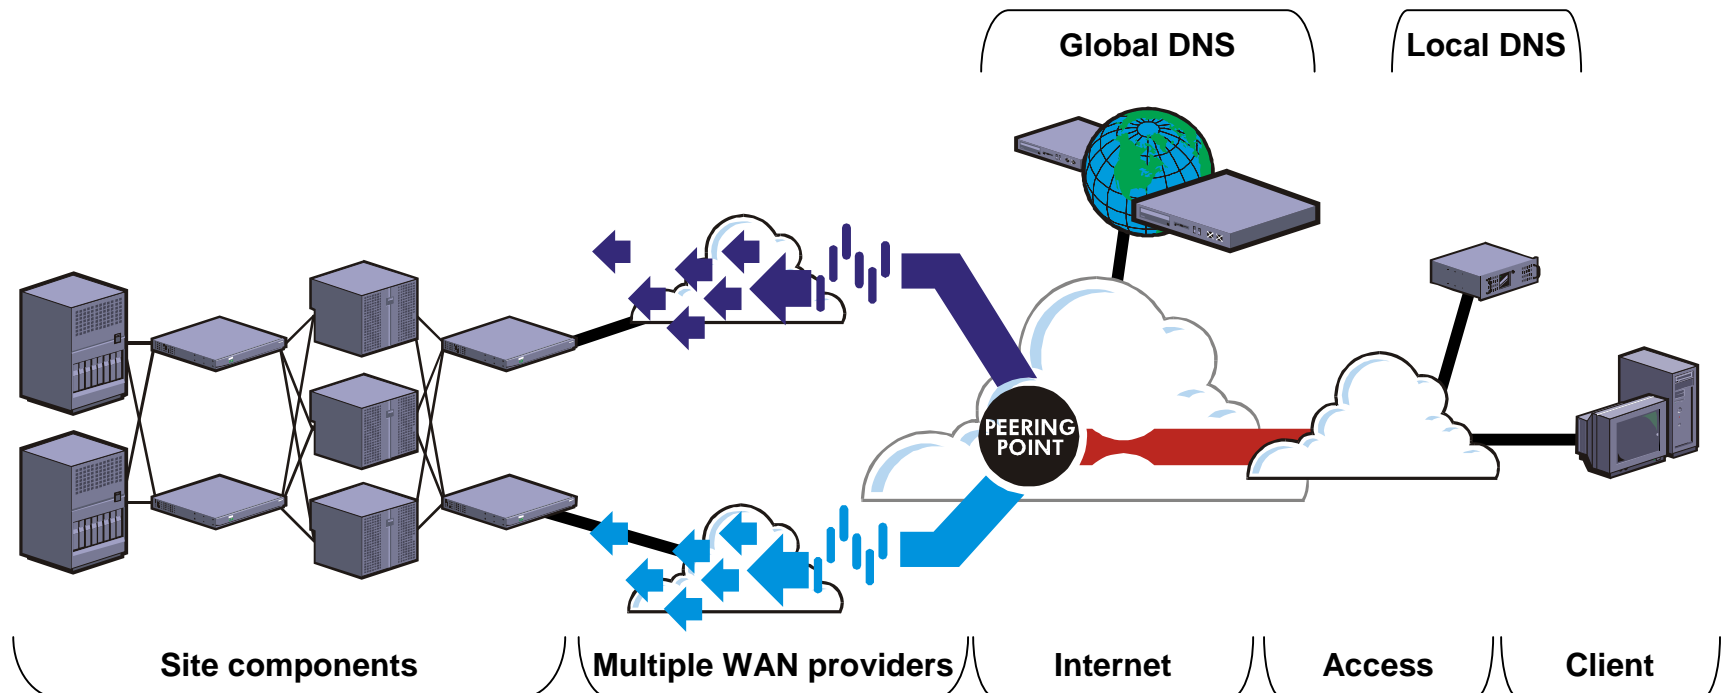
- Technical challenges
 - What makes sense to distribute?
 - What’s the “best” point?
 - How do we force clients there without breaking things?

How a CDN works

- A look at site architecture
- Picking a fulfillment point
- Directing the client there
- Publishing content to the edge


A look at site architecture

- To understand CDNs, we need to understand how sites get content
- There are many components in an online system




A look at site architecture

- Let's simplify things a bit...




Step one (oversimplified): pick a fulfillment point




- Client resolves the IP address of the site
 - www.mysite.com

Step one: pick a fulfillment point




- Client requests an object from server
 - Index.html

Step one: pick a fulfillment point




- CDN client decides which cache is logically closest to the client
 - Based on location, BGP, content, etc.

Step two: direct the client there



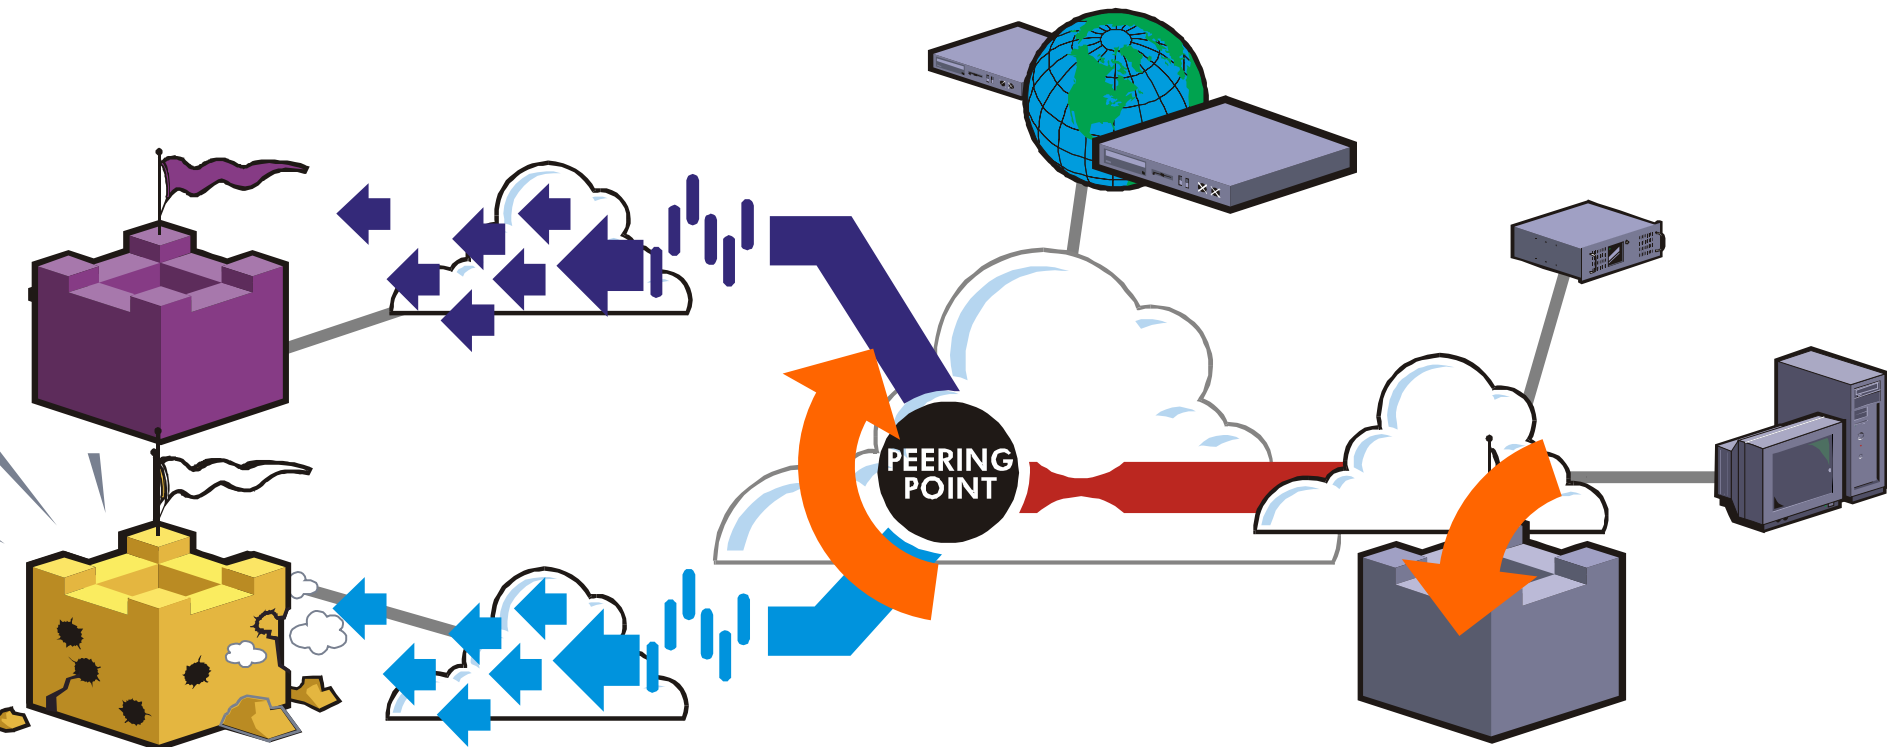
- Server responds with a custom-built page
 - Links in the page point to the newly chosen location, not the original site

Step two: direct the client there




- Client retrieves subsequent content from local cache
 - GIFs, etc.

Handling outages




- If one of the data centers is unavailable
 - Send the client elsewhere
 - Send the client to a cached copy

Publishing the content there reactively



- Can be demand-based
 - First time it's requested, it's stored

Publishing the content there proactively



- Can be proactive
 - When large loads are expected


Early content delivery systems were simply smart DNS and caching

- Primary benefits were performance gains by not crossing peering points
- Priced around \$2,000 per megabyte
- First hit always went to the root site
- Required the site agent to generate custom links

Now the CDN owns the root for your domain

- Client hits a “core” cache server within the hierarchy
- CDN system does triage to better refine things
- Main benefits are reduced web server footprint, better availability

A slightly less simple view



- CDN owns the DNS for the client
 - Even the “first hit” to the site is handled by a cache point

Next step is dynamic processing at the edge

- Topical subsections of site with some “thinking”
- Permissions, security, freshness are what matter
- Main benefits are distributed processing, reduced application server footprint

Emerging technologies

- New clients
- Edge-side includes
- Hardware-based

Emerging technologies: New clients

- PDA and wireless devices have different formats
- Edge presentation can simplify the core



Emerging technologies: ESI

- Allows the inclusion of multiple “fragments” into a page
- Assembled at the edge
- Can handle variables
 - Either to perform calculations or for insertion in HTML output
- Can do computation and handle errors
- This is relevant because **it changes where things get processed**
 - For more info, www.esi.org

Geekfest: ESI Inclusion



- Edge composes pages by assembling included content
- each “fragment” has own cacheability & handling attributes

```
<esi:include src="http://example.com/1.html"  
  alt="http://bak.example.com/2.html"  
  onerror="continue"/>
```

Geekfest: ESI Variable support



- Uses variables based on HTTP request attributes
- Can be interpreted or written into processed page

```
<esi:vars>
  
</esi:vars>
```


Geekfest: ESI Conditional processing



- Conditional logic influences how a template is processed

```
<esi:choose>
  <esi:when
    test="$(HTTP_COOKIE{group})=='Advanced'">
    <esi:include
      src="http://www.example.com/advanced.html"
    />
  </esi:when>
  <esi:when
    test="$(HTTP_COOKIE{group})=='Basic
    User'">
    <esi:include
      src="http://www.example.com/basic.html"/>
  </esi:when>
  <esi:otherwise>
    <esi:include
      src="http://www.example.com/new_user.html"
    />
  </esi:otherwise>
</esi:choose>
```

Geekfest: ESI Exception/error handling

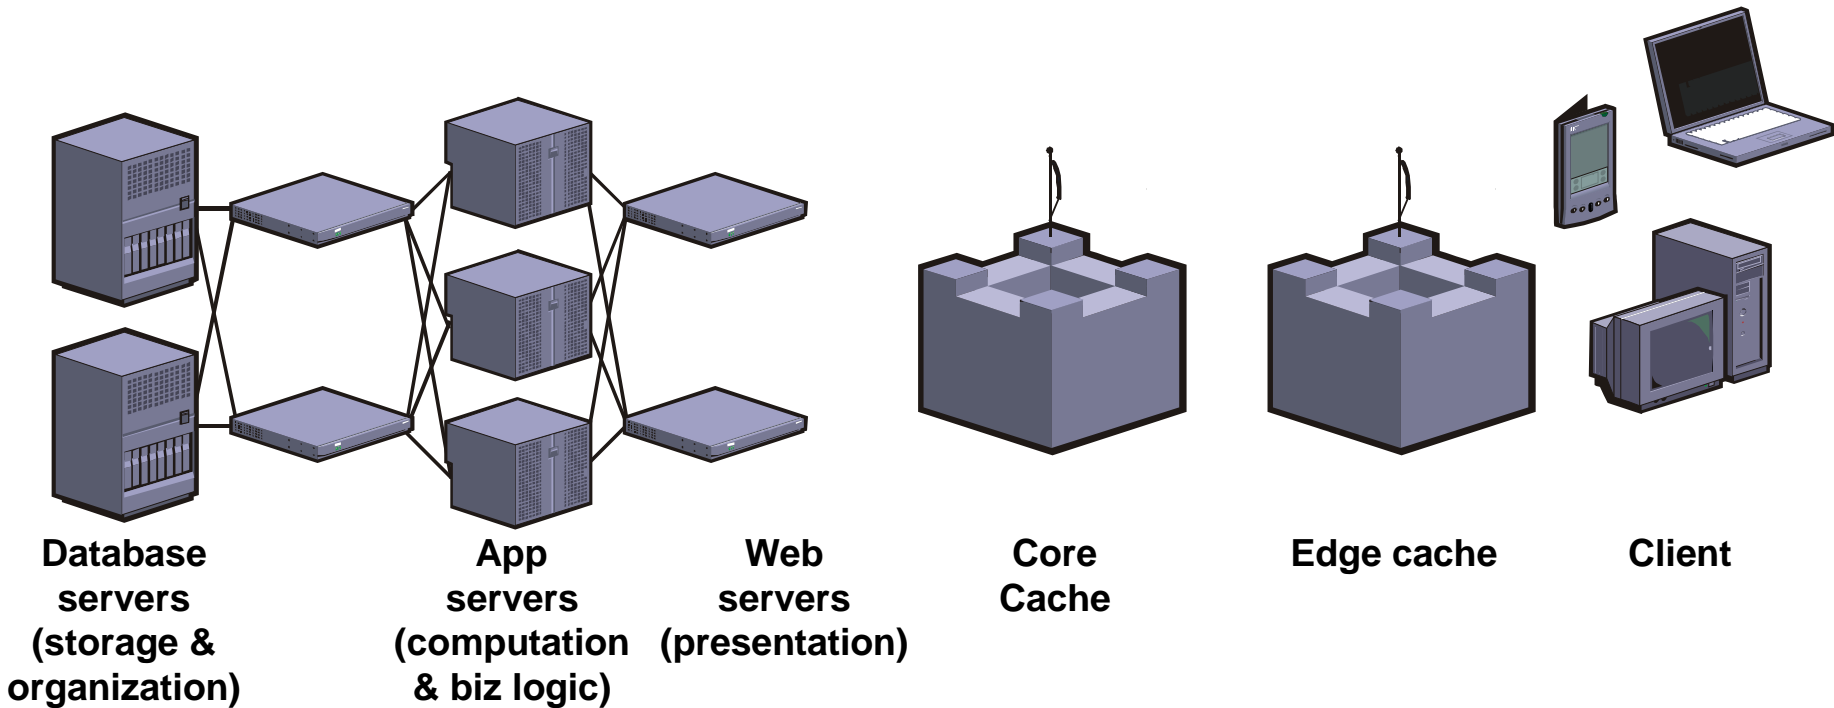


- specification of alternate and default resources

```
<esi:try>
  <esi:attempt>
    <esi:comment text="Include an ad"/>
    <esi:include
src="http://www.example.com/ad1.html"/>
  </esi:attempt>
  <esi:except>
    <esi:comment text="Just write some
HTML instead"/>
    <a
href=www.akamai.com>www.example.com</a>
  </esi:except>
</esi:try>
```

Different levels of distribution

- Architecture starts by deciding what to distribute



One-time content

Quasi-dynamic content

Static content

Custom format

Using a CDN

- Increasingly easy to use
 - Now like buying bandwidth
- Reduces control somewhat
 - Lack of ability to steer traffic, own DNS
 - The idea of “closer” varies
 - Router hops, latency, site load, packet loss, or other factors
 - Need to keep content fresh
- Possible security, copyright issues
 - Storage of copyrighted materials
- As processing gets distributed, transaction control becomes critical
 - The “last ticket” problem

Alternatives to a CDN

- Multi-homed networks
 - Internap
- Building your own client-side caches
- Keeping your sites lean

The business-enabled CDN

- Has a good push model
 - Proactively, for content with expected popularity
 - Reactively, when demanded
- Has a broad footprint domestically/internationally
- Is reasonably priced
- Has strong technology
 - Breadth of cache distribution
 - Algorithm effectiveness
 - Overhead on clients
 - Robust security and privacy of content or logic
 - Ability to tune and signal the network
- Has a compelling roadmap
 - New clients, edge processing
- Has a good rate of innovation
- Is easy to operate and scale up/down on demand

Shopping list

- “Name brand” CDN vendors
 - Akamai
 - Adero
 - Inktomi
 - Mirror Image
- Industry associations
 - Contentbridge
- Build-your-own caching vendors
 - Cacheflow
 - Infolibria
 - Inktomi